

EXHIBIT 18

This document describes the suggested cli changes to accommodate IPv6.

Currently the following features are supported:

- o ICMPv6, Neighbor Discovery
- o Transport over: Ethernet, PPP, cisco HDLC, FDDI, Frame Relay, ATM
- o Transition mechanisms: configured tunneling, automatic tunneling, 6to4 and 6over4.
- o DNS, - only AAAA records.
- o ping, trace, telnet
- o file copying commands (tftp/ftp), net booting over tftp
- o BGP4+, RIPng
- o simple access-lists, prefix-lists, route-maps.

Target IOS version: 12.1(5)T

Global configuration commands:

```
ipv6 access-list <name> permit [<prefix> [*] any] {<prefix> [*] any}
ipv6 access-list <name> deny [<prefix> [*] any] {<prefix> [*] any}
no ipv6 access-list <name>
IPv6 access lists are identified by user selected names. Access lists
are defined by a list of "permit" and "deny" statements. These
statements each specify two search keys, a first order key which
is mandatory and entered in last and an optional second order (which
defaults to 'any' if omitted).
```

Each key specifies a prefix or class or prefixes to match, depending on the "*" flag. Wild-card keys will match any prefix as or more significant than the key that matches the leading portion of the key prefix. The keyword 'any' is equivalent to '0::0/0 *'.

[no] ipv6 hop-limit <value>
 Configure the system to use <value> as the IPv6 Hop Limit value used in Router Advertisements and in all IPv6 packets generated within the system. The default value is 64.

[no] ipv6 icmp error-interval <milliseconds>
 Rate limits the number of error ICMPv6 messages sent.
 Default is: maximum one every 500ms.

[no] ipv6 prefix-list sequence-number
 [no] ipv6 prefix-list <name> description <text>
 [no] ipv6 prefix-list <name> [seq <seq-value>] permit|deny <ipv6prefix>
 {{ge <min-value>} [le <max-value>]}

 no ipv6 prefix-list <name>
 Mimics the 'ip prefix-list' commands.

[no] ipv6 route <prefix> {<next-hop> | <interface>} [<distance>]
 Configures a static IPv6 prefix route. <prefix> could also be a host name when configuring static host routes. <next-hop> is the host name or IPv6 address of the next-hop to reach the destination prefix. The next-hop address need not be directly connected. Recursion is done to find the real physical next-hop. When <next-hop> is not used, <interface> can be used to direct static routes out point-to-point interfaces like serial links or tunnels. The default value for <distance> is 1 which gives static routes precedence over any other

type of route with the exception to Connected routes.

[no] ipv6 router {<protocol>} [<protocol specific>]
 Enable an IPv6 routing process. (Does not apply for
 IPv6 BGP configuration, which is under 'router bgp'.)

[no] ipv6 unicast-routing
 Enables forwarding of IPv6 unicast datagrams. The default setting
 is disabled.

Modified commands:

[no] ip host <name> [<port>] <ipv6addr>|<ipv4addr> [<ipv6addr>|<ipv4addr> ...]
 Configures a static domain name for the specified address. Up to
 four addresses are allowed per domain name.

[no] ip name-server <ipv6addr|ipv4addr> [<ipv6addr|ipv4addr> ...]
 This command now accepts any combination of IPv4 or IPv6
 addresses, up to a total of six. DNS queries will be sent
 using the IP layer implied by the name-server server, that is,
 if you specify an IPv6 address, the DNS request will use IP
 version 6 datagrams.

[no] dialer-list <no> protocol ipv6 {permit | deny | list <acl-name>}
 Added support for IPv6 acls.

[no] logging {<hostname|ipv6addr|ipv4addr>}
 [no] tacacs-server host {<hostname|ipv6addr|ipv4addr>}
 [no] radius-server host {<hostname|ipv6addr|ipv4addr>}
 Added support for IPv6 addresses and host names.

Interface configuration commands:

[no] ipv6 address <ipv6addr>[/<prefix-length>] [<link-local>]
 Configures an IPv6 address (and enables IPv6) on the interface. When
 the keyword "link-local" is specified, the link-local unicast address
 will override the one selected by the system, which is the IEEE address
 for the interface. When <ipv6addr> is specified, it must be in
 the form documented in RFC2373 where the address is specified in
 hexadecimal using 16-bit values between colons. Optionally, a double
 colon may be used when consecutive 16-bit values are denoted as zero.
 In addition, optionally the last 4 bytes may be specified in IPv4
 dotted decimal. You can configure multiple IPv6 addresses per
 interface but only one link-local address. To delete all addresses,
 use the no form without any parameters.

[no] ipv6 address <ipv6prefix>/<prefix-length> eui-64
 Configures an IPv6 address (and enables IPv6) on the interface
 using a EUI-64 style "Interface ID" in the low order 64 bits.
 (If the <prefix-length> specified is greater than 64, the
 prefix bits will have precedence over the Interface ID.)

[no] ipv6 enable
 Enables IPv6 on an interface with no address configured. Default

value is disabled. Will auto-configure an IPv6 link-local unicast address on the interface as part of enabling the interface.

The "no ipv6 enable" command will not disable IPv6 on an interface if there are explicitly configured addresses on the interface.

[no] **ipv6 mtu <bytes>**

Configures the Maximum Transmission Unit for IPv6 datagrams on an interface. The default value is the link MTU. If a non-default value is configured, an MTU option will be included in Router Advertisements.

[no] **ipv6 nd ns-interval <milliseconds>**

Configures the interval between IPv6 Neighbor Solicit transmissions. This value will be included in all IPv6 Router Advertisements sent out this interface. Very short time intervals are not recommended in normal IPv6 operation. By default, the value 0 (unspecified) is advertised in Router Advertisements and the value 1000 is used for the router's own Neighbor Discovery activity. When a non-default value is configured, it will be both advertised and used by the router itself.

[no] **ipv6 nd ra-interval <seconds>**

Configures the interval between IPv6 Router Advertisement transmissions sent out this interface. The value for this should be less than or equal to the IPv6 Router Lifetime if this is a default router. The default is 200 seconds. To prevent synchronization with other IPv6 nodes the actual value used may be randomly adjusted to within +/- 20% of the specified value.

[no] **ipv6 nd ra-lifetime <seconds>**

Configures the lifetime of a Router Advertisement. This value will be included in all IPv6 Router Advertisements sent out this interface. If the router is NOT a default router, this will have value 0. If the router is a default router, this value will be non-zero and should not be less than the minimum Router Advertisement interval. The default value is 1800 seconds.

[no] **ipv6 nd reachable-time <milliseconds>**

Configures the time that a remote node is assumed reachable after some reachability confirmation event has occurred. The value for this is included in all IPv6 Router Advertisements sent out this interface. This is used to help detect dead neighbors. Shorter values discover dead neighbors more quickly at the expense of more bandwidth consumed and more processing costs in all IPv6 systems. Very short time intervals are not recommended in normal IPv6 operation. By default, the value 0 (unspecified) is advertised in Router Advertisements and the value 1,800,000 (30 minutes) is used for the router's own Neighbor Discovery activity. When a non-default value is configured, it will be both advertised and used by the router itself.

[no] **ipv6 nd suppress-ra**

Control transmission of IPv6 Router Advertisements on the interface. The default is to send Router Advertisements (RA's) on Ethernet, FDDI, or Token Ring interfaces if ipv6 unicast-routing is enabled. Use the command "ipv6 nd suppress-ra" to turn off RA's on LAN

interfaces. On other types of interfaces, the default is to never send an RA. Use the command "no ipv6 nd suppress-ra" to send RA's on interfaces such as serial or tunnel interfaces.

[no] ipv6 nd prefix-advertisement <routing-prefix>/<length> <valid-lifetime>
 <preferred-lifetime> [onlink | autoconfig]

Explicitly configure which IPv6 routing prefixes are advertised in an Router Advertisement. If this command is not used, then the router will advertise all prefixes configured on the interface originating the Router Advertisements.

[no] ipv6 nd managed-config-flag

Defaults is OFF. When OFF, Router Advertisements sent out this interface have the "Managed Address Configuration Flag" turned off. Hosts are thus permitted to use IPv6 stateless auto configuration to create global unicast addresses for themselves.

[no] ipv6 nd other-config-flag

This defaults to OFF. When OFF, Router Advertisements sent out this interface have the "Other Stateful Configuration Flag" turned off meaning that hosts should not use DHCP to perform configuration of "other" (meaning non-address) parameters.

[no] ipv6 traffic-filter <acl-name> {in | out}

Filter incoming (outgoing) traffic according to the specified access list.

[no] ipv6 unnumbered <interface>

Configures an interface to enable IPv6 without requiring a global address. <interface> is an interface that does have a global address configured using the "ipv6 address" command. This command is used to reduce address administration for a network administrator. Any packets that are originated over an unnumbered interface will use the first configured global address, of the numbered interface it points to, as the source address in the IPv6 header.

This command can only be used on point to point interfaces.

Tunnel interfaces:

[no] tunnel mode ipv6ip

Configures a static tunnel interface to encapsulate IPv6 packets in IPv4. This interface can be used like any other interface (static routes can point to it or a dynamic routing protocol can run over it). This is "configured tunneling" according to RFC 1933.

[no] tunnel mode ipv6ip auto-tunnel

Configures IPv6 in IPv4 automatic tunneling. Automatic tunneling is performed when a destination address is an IPv6 packet containing an IPv4 compatible address as specified in RFC2373.

[no] tunnel mode ipv6ip 6to4

This implements the 6to4 transition mechanism as defined in draft-ietf-ngtrans-6to4-03.txt. The tunnel source command must point to an interface with an IPv4 address configured. This

IPv4 address will be used as the NLA in the 6to4 prefix.

[no] tunnel mode ipv6ip 6over4

This implements the 6over4 transition mechanism as defined in RFC2529. The tunnel source command has to point to the interface attached to the IPv4 multicast domain. 239.192/16 is used as the default IPv4 multicast block. Currently this can not be overridden. Neighbor Discovery is automatically enabled on this interface.

ATM interface commands:

[no] ipv6 <address> atm-vc <vcid>

Here is an example configuration:

int atm1/0

atm pvc 1 <VPI> <VCI> aal5snap
map-group foo

map-list foo

ipv6 5F00:6D00::5 atm-vc 1

Rip interface commands:

[no] ipv6 rip <tag> enable

Configures RIP routing on the interface.

If you wish to run RIP on an unnumbered tunnel interface you must enable RIP routing on both the tunnel interface and on the physical interface the tunnel interface is using as the source of its IP address.

[no] ipv6 rip <tag> default-information originate

Originate the default route (::/0) and include in updates sent from this interface.

[no] ipv6 rip <tag> default-information only

Originate the default route (::/0). Suppress sending any routes except the default route on this interface.

[no] ipv6 rip <tag> summary-address <prefix>/<length>

A means of compressing routing information. If the first <length> bits of a route match the given prefix, the prefix will be advertised instead. Multiple routes are thus replaced by a single route whose metric is the lowest metric of the multiple routes. This command may be used multiple times.

Router submode configuration commands:

in submode: [no] ipv6 router rip <process-name>

[no] distance <distance>

Set the administrative distance for this RIP process. The default value is 120. If two RIP processes attempt to insert the same route into the same routing table, the one with the lower administrative distance will

take precedence.

[no] distribute-list prefix-list <pfx-name> {in|out} [<interface>]
Apply a simple access list to RIP routing updates received or sent on an interface. If no interface is specified, the access list will be applied to all interfaces.

exit

Exit from IPv6 RIP configuration mode

[no] metric-offset <number>
Normally a RIP metric is incremented by one before being entered in the routing table. This command changes that increment to a new value between 1 and 16.

[no] poison-reverse

Perform poison reverse processing of updates, that is, when we advertise networks on interfaces from which we learned them, advertise an unreachable metric. Off by default. If both split horizon and poison reverse are enabled, only split-horizon processing will be done.

[no] port <udp-port> multicast-group <multicast-address>

Configure the RIP routing process to use the specified UDP port and multicast address. If this command is not given, the standard port 521 and multicast address FF02::9 will be used. If two RIP processes are assigned the same UDP port, the second will not startup until the configuration conflict is resolved. Two RIP processes may use the same multicast address.

[no] redistribute {connected|static|bgp|rip} [<rp-specific>] [route-map <name>]

Advertise static routes as if they were directly connected. This command should be used with caution because advertising routes that are not directly connected can cause routing loops if improperly used.

[no] split-horizon

Perform split horizon processing of updates, that is, suppress advertising networks on interfaces from which we learned them. On by default.

[no] timers <update> <expire> <holddown> <garbage-collect>

Configures RIP routing timers for the RIP process. Periodic updates are sent every <update> seconds. Networks not heard about in <expires> seconds are marked unreachable. Information about unreachable networks is ignored for a further <holddown> seconds. If you do not want hold-downs, use 0 seconds. An expired entry is deleted in <garbage-collect> seconds after either expiration or hold-down termination. The default values are

timers 30 180 180 120

BGP:

We are using the new multiprotocol BGP for configuring IPv6 BGP. All relevant commands accepting either prefixes or a neighbor address has been extended to accept a IPv6 address/prefix in

the IPv6 address family.

See 'IOS CLI for BGP/MPLS VPN', ENG-29568 for detail on the 'address-family' command.

e.g:

```
[no] router bgp <AS>
  neighbor <ipv6address> <commands>
  address-family ipv6
    network <ipv6prefix>
    aggregate-address <ipv6prefix>
  address-family ipv4
  [...]
```

Route-map configuration commands

```
-----
```

[no] match ipv6 address <access/pfx-list>
 [no] match ipv6 next-hop <access/pfx-list>
 [no] match ipv6 route-source <access/pfx-list>
 Use the given access-list to match the address, next-hop and route-source fields of routes. The specific interpretation of those fields depends on the routing protocol.

Line configuration commands

```
-----
```

[no] ipv6 access-class <acl> {in|out}
 Filter connections based on an IPv6 access list.

=====
 Debug commands:

[un]debug ipv6 packet
 Enables IPv6 packet-level debugging. Default setting is off.

[un]debug ipv6 icmp
 Enables ICMPv6 debugging, except Neighbor Discovery.
 Default setting is off.

[un]debug ipv6 nd
 Enables ICMPv6 Neighbor Discovery debugging. Default setting is off.

[un]debug ipv6 routing
 Enables IPv6 routing table event activity. Default setting is off.

[un]debug ipv6 rip
 Enables RIP routing protocol debugging, such as when RIP packets are sent/received on an interface. Default setting is off. Use of this command on a busy network might seriously impact router performance.

[un]debug ipv6 rip <interface>
 Same as "debug ipv6 rip" except that packet input and output events are restricted to the specified interface.
 Default setting is off. Use of this command on a busy network will seriously impact router performance.

Clear commands:

clear ipv6 neighbors

Deletes all entries in the IPv6 ND cache.

clear ipv6 prefix-list [<name>] [<ipv6prefix>]

clear ipv6 route <ipv6addr-or-name> | *

Clears the IPv6 routing table. If <ipv6addr-or-name> is specified, only a single route is deleted. When "*" is specified, the entire routing table is deleted.

When "*" is specified, the per-destination MTU cache is also cleared.

clear ipv6 traffic

Reset the counters in the "show ipv6 traffic" display.

Show commands:

show ipv6 access-list

Display the currently defined access lists.

show ipv6 interface [<interface>]

Displays IPv6 interface related parameters and addresses.

show ipv6 mtu

Displays the per-destination MTU cache.

show ipv6 neighbors [<ipv6addr-or-name> | <interface>]

Displays neighbor adjacency entries from the IPv6 Neighbor Discovery (ND) table. Includes the state of the adjacency entry, its lifetime, and the associated MAC and IPv6 addresses. If an interface or an IPv6 address or an IPv6 name is specified, then only the related entries will be displayed.

show ipv6 prefix-list [summary|detail] [<name>]

show ipv6 prefix-list <name> seq <seq-num>

show ipv6 prefix-list <name> <ipv6prefix> [first-match | longer]

Mimics the ip prefix-list show commands.

show ipv6 rip

Display status of the various RIP processes.

show ipv6 route [connected | local | static] | [<prefix> | <addr-or-name>]

Displays the IPv6 routing table. When <prefix> or <addr-or-name> is supplied, only a single route is displayed. When <prefix> or <addr-or-name> is supplied, a longest match lookup is performed from the routing table. When one of the keywords "connected", "local", or "static" is supplied, only those type of routes are displayed.

show ipv6 route summary

Displays number of routes per route source. Displays number of routes per prefix length.

show ipv6 routers

Displays Router Advertisement (RA) information received from

on-link routers. Routers advertising RA parameters differing from those configured for the interface are flagged as "conflicting".

show ip sockets

Enhanced to display IPv6 related socket information in addition to the previous IPv4 information.

show ipv6 traffic

Displays IPv6 related traffic statistics

show ipv6 tunnel

For each tunnel running IPv6 display the tunnel unit number, the name of the dynamic routing protocol in use, the time of last input, the number of packets input, and the description string as set by the "description" interface subcommand.

Exec commands:

traceroute [ipv6] <destination>

Traces the route for IPv6 packets between this node and the destination address. Reports on the time to each hop and, if the remote end responds, reports the identity of each hop.

ping [ipv6] <destination>

Sends ICMPv6 echo request packets to <destination>. <destination> can be an IPv6 host name or address. In enable mode, you can be prompted interactively for ping options.

telnet <hostname>|<ipv4addr>|<ipv6addr> [/ipv4 | /ipv6]

The switches /ipv4 and /ipv6 will force the use of the appropriate network layer, when DNS resolves a hostname with both IPv6 and IPv4 addresses. If no switches are specified, we will first try IPv4, then IPv6 to make the connection.

TFTP file loading

All interactive commands that download or upload TFTP files now accept IPv6 numeric addresses or host names.
For hostnames which returns both A and AAAA/A6 records, we will first try the A record, if that doesn't succeed, we will try the AAAA/A6 record(s).

DNS

The usual name -> address and address -> name lookups are supported.

Rommon commands

b now accepts an IPv6 address.